# Technical Note: Context Specialization Metric (CSM)

## Motivation

Context-conditioned neural architectures (e.g. Atlas neural networks) associate different local linear mappings with different context values. A natural question is:

> **How informative is a particular context variable?**

Rather than measuring the context directly, the **Context Specialization Metric (CSM)** measures how much the learned computations diverge after training.

---

## Core Idea

Consider an Atlas layer conditioned on a context variable (e.g. time of day).

Each context value selects a local linear mapping

- $W_1$
- $W_2$
- ...
- $W_n$

After backpropagation training, compare the learned mappings.

The hypothesis is simple:

- **Nearly identical mappings** → Context is computationally uninformative.
- **Distinct mappings** → Context is computationally informative.
- **Clusters of mappings** → Context contains a smaller number of computationally distinct modes.

The specialization emerges automatically from optimization rather than being manually specified.

---

## Possible Distance Measures

Compare mappings using metrics such as:

- Frobenius norm.
- Spectral norm.
- Cosine similarity.
- Principal angles between dominant subspaces.
- Output differences over representative datasets.

The CSM may be computed as an average pairwise distance or another suitable specialization statistic.

## Interpretation

### Low CSM

- Context ignored.
- Similar computations for all context values.
- Context variable may be unnecessary.

### High CSM

- Strong specialization.
- Different computations learned for different contexts.
- Context contributes significantly to performance.

### Clustered CSM

The learned mappings naturally form groups.

Example:

- Morning ≈ Dawn
- Afternoon ≈ Evening
- Night distinct

This suggests that the supplied context contains fewer computationally meaningful states than originally assumed.

---

## Layer-Dependent Context

Different layers may specialize differently.

For example:

Early perception

- Spatial context important.
- Time context ignored.

Planning

- Mission context important.
- Temporal context important.

Motor control

- Behavioral state important.
- Circadian context weak.

The CSM therefore provides a profile of **where** each context variable influences computation.

## Applications

Potential context variables include:

- Behavioral state.
- Finite automata state.
- Time of day.
- Circadian phase.
- Navigation state.
- Location.
- Goal state.
- Safety mode.
- Battery level.
- Physiological state.
- Atlas geometric context.

---

## Model Simplification

If multiple mappings converge to nearly identical solutions:

- Merge equivalent mappings.
- Compress the context representation.
- Reduce parameter count.
- Discover computational equivalence classes automatically.

The CSM therefore provides both an analysis tool and a possible model compression strategy.

---

## Research Questions

- Which context variables produce the greatest specialization?
- Which layers benefit from explicit context?
- How many computational modes are actually required?
- Can context variables be automatically merged after training?
- Does higher specialization correlate with improved generalization or sample efficiency?

---

## Summary

The **Context Specialization Metric (CSM)** measures the extent to which explicit context causes local computations to specialize after learning. Rather than asking whether a context variable exists, the CSM asks whether it changes the computation enough to justify dedicated parameters. This provides a simple quantitative framework for evaluating explicit context in context-conditioned neural architectures.